TABLE OF CONTENTS

SECTION **PAGE**

5.10 Data Storage Controller 2
5.10.1 Introduction 2
5.10.2 Storage System Requirements 2
5.10.3 Storage Protocol Requirements 3
5.10.4 NAS Interface Requirements 5
5.10.5 SAN Interface Requirements 6
5.10.6 CNA Interface Requirements 6
5.10.7 IP Networking Requirements 6
5.10.8 Name Services Requirements 7
5.10.9 Security Services Requirements 8
5.10.10 Interoperability Requirements 9
5.10.11 Class of Service & Quality of Service Requirements 9
5.10.12 Virtualization Requirements 10

5.10 Data Storage Controller

5.10.1 Introduction

A Data Storage Controller (DSC) is a specialized multi-protocol computer system with an attached disk array that together serves in the role of a disk array controller and end-node in B/P/C/S networks. A DSC provides data storage and services as depicted below in Table 1.

Storage Type	Network Infrastructure
Network Attached Storage (NAS)	• Ethernet • Multi-protocol: NFS, CIFS, iSCSI, HTTP, FTP
Storage Array Network (SAN)	• Fibre Channel (FC) • Fibre Channel Protocol (FCP)
Converged Network	• Ethernet • Multi-protocol: NFS, CIFS, iSCSI, HTTP, FTP • Fibre Channel over Ethernet (FCoE) • Data Center Bridging (DCB)

Table 1 – DSC Data Storage and Service Types

DSC features and capabilities listed in this section may be offered as a part of a unified capability offering associated with other products on the APL. The definitions for Data Storage Controller are found in Section A1, Definitions, Abbreviations, Acronyms and References.

5.10.2 Storage System Requirements

1. **[Required: DSC]** The system shall provide a Redundant Array of Independent Disks (RAID) for multiple disk drives. The system shall provide a configuration option to select the specific RAID level to be provisioned in the disk array. The RAID levels available for use shall be subject to the specific vendor implementation. The default level shall be dual parity RAID-6 for SATA drives and RAID-5 for Serial Attached Small Computer Systems Interface (SCSI) (SAS) and FIBRE Channel Drives.

2. **[Required: DSC]** The system shall be capable of 99.9% availability in High Availability (HA) mode. The system shall have a configurable parameter that allows this mode to be enabled and disabled.

3. **[Required: DSC]** The system shall provide a management control function for low-level system monitoring and control functions, interface functions and remote management. The management control function shall provide an Ethernet physical interface(s) and status. The monitoring shall include initial system check, system cooling fans, temperatures, power supplies, voltages and system power state tracking and logging.

4. **[Required: DSC]** The system shall provide data storage replication (e.g. mirroring) services (IPv4 and IPv6) between systems that are configured as source & destination replication pairs.

The replication operations shall provide capabilities for data backup replication, system replication & migration and system Disaster Recovery (DR) services in support of Continuity of Operations Planning (COOP).

- **[Required: DSC]** When the system interfaces to a Integrated Data Protection (IDP) service and the IDP makes copies of data storage information on to another DSC for the purposes of periodic data storage backup, DR/COOP, migration and data archiving operation the system replication service shall complete the replication regardless of the host connection protocols used between the application servers and the DSC.
- **[Required: DSC]** The system Replication & Migration services shall provide capabilities to replicate data storage and configuration information on to another stand-by DSC system for the purposes of migrating data storage information.
- **[Required: DSC]** The system Disaster Recovery services shall provide capabilities to replicate data storage and configuration information on to another stand-by DSC system for the purposes of DR/COOP.

NOTE: The approach described above provides the threshold capability. Other replication techniques are permitted to ensure communication optimization.

5. **[Conditional: DSC]** The system shall provide configurable modes for replication (mirroring) operations between the source DSC and the destination DSC. During replication, both the source and the destination must be in a known good state. The configurable modes shall be Asynchronous or Synchronous and are depicted below in Table 2.

Replication Mode	Description
Asynchronous (Async)	Incremental, block-based replication between DSCs that occurs as frequently as once per minute by scheduling or manually entering a command to trigger the replication operations.
Synchronous (Sync)	Real-time replication between DSCs that occurs as data is stored or as it changes.

Table 2 – Replication Operation Modes.

5.10.3 Storage Protocol Requirements

1. **[Required: DSC]** The system shall provide a Network File System version 3 (NFSv3) server for file systems data I/O.

2. **[Conditional: DSC]** The system shall provide a Network File System version 4 (NFSv4) server for file systems data I/O.

3. **[Conditional: DSC]** The system shall provide a Network File System version 4.1 (NFSv4.1) server, including support for parallel NFS (pNFS) for file systems data I/O.

4. **[Required: DCS]** The system shall provide Common Internet File System version 1.0 (CIFSv1.0) server for file systems data I/O.

5. **[Conditional: DCS]** The system shall provide Common Internet File System version 2.0 (CIFSv2.0) server for file systems data I/O.

6. **[Conditional: DCS]** The system shall provide Internet Small Computer System Interface (iSCSI) server (target) operations for data I/O of Logical Units (LUNs) to clients (initiators).

7. **[Conditional: DCS]** The system shall provide Fibre Channel Protocol server (target) operations for data IO of Fibre Channel protocol Logical Units (LUNs) to clients (initiators).

8. **[Conditional: DCS]** The system shall provide Fibre Channel over Ethernet (FCoE) server (target) operations for data I/O of Fibre Channel protocol Logical Units (LUNs) to clients (initiators).

9. **[Conditional: DCS]** The system shall provide a Hyper Text Transfer Protocol Secure (HTTPS) server for file system data I/O and management access to the storage controller operating system. The session shall be secured with Secure Sockets Layer (SSL) or Transport Layer Security (TLS) as per IETF RFC 5246 and shall comply with Section 5.4 requirements for that protocol.

10. **[Required: DCS]** The system shall provide Secure Shell version 2 (SSHv2) or Secure Sockets Layer (SSL) for management access to the storage controller operating system. The SSHv2 or SSL implementation shall comply with Section 5.4 requirements for that protocol.

11. **[Conditional: DCS]** The system shall provide Web-based Distributed Authoring and Versioning (WebDAV) as per IETF RFC 4918, in support of Cloud-based virtualized storage infrastructures.

12. **[Conditional: DCS]** The system shall implement the Representational State Transfer (REST) software architecture for distributed hypermedia systems and Cloud-based virtualized storage infrastructures.

13. **[Conditional: DCS]** The system shall implement the SNIA Cloud Data Management Interface (CDMI) standard.

14. **[Required: DCS]** The system shall provide Global Name Space (GNS) or Single Name Space functionality. The GNS functionality shall provide the capability to aggregate disparate and remote network based file systems to provide a consolidated view to reduce complexities of localized file management and administration. The GNS functionality shall provide large (i.e., 14 PB or greater) working pools of disks, transparent data migration, and serve to reduce the number of storage mount points and shares. Each system shall have a dedicated and unique GNS.

NOTE: A global name space functionality is provided with the assumption that it will only be used in deployments where latency is below 200 ms.

5.10.4 NAS Interface Requirements

1. **[Required: DSC]** The system shall provide physical interfaces for Gigabit Ethernet (GE) and 10 Gigabit Ethernet (10 GE) services in conformance with IEEE 802.3 for Ethernet LAN interfaces.

2. **[Required: DSC]** The system shall be able to provision, monitor, detect faults, and restore Ethernet services in an automated fashion.

3. **[Required: DSC]** The system shall provide physical interfaces for Out-of-Band (OOB) management access and services with 10/100 Mbps Ethernet interfaces as a minimum. Services shall include remote access with at least one of the following protocols: Secure Shell version 2 (SSHv2), SSL, HTTPS, and SNMPv3; and the protocols shall be secured in accordance with Section 5.4.

4. **[Required: DSC]** When the system uses Ethernet, Fast Ethernet, Gigabit Ethernet and 10 Gigabit Ethernet interfaces, the interfaces shall be auto-sensing, auto-detecting and auto-configuring with incoming and corresponding Ethernet link negotiation signals.

5. **[Required: DSC]** Ethernet services of the system and the Logical Link Inter-Working Function (IWF) of the system shall terminate the Media Access Control (MAC) layer of Ethernet as described in Ethernet Standard IEE 802.3.

6. **[Required: DSC]** Ethernet services of the system shall support jumbo frames with a configurable Maximum Transmission Unit (MTU) of 9000 bytes or greater, excluding Ethernet encapsulation. When Ethernet encapsulation is included in the frame size calculation an additional 22 bytes must be included for the MAC header (14 bytes), the VLAN tag (4 bytes) and the CRC Checksum (4 bytes) fields in the Ethernet frame, resulting in a maximum of 9022 bytes or greater. The system shall also support a configurable MTU between 1280 bytes and 1540 bytes to ensure packets can transit type 1 encryptors. The system default MTU shall be 1540 bytes.

7. **[Required: DSC]** Ethernet services of the system shall provide Virtual LANs (VLANs) as per IEEE 802.1Q and shall allocate a unique Ethernet MAC address to each VLAN.

8. **[Required: DSC]** Ethernet services of the system shall support “Link Aggregation” as per IEEE 802.3ad or IEEE 802.1AX-2008 and use with Link Aggregation Control Protocol (LACP).

9. **[Required: DSC]** Ethernet services of the system shall provide Link Layer Discovery Protocol (LLDP) as per IEEE 802.1AB.

5.10.5 SAN Interface Requirements

1. **[Conditional: DSC]** The system shall provide Fibre Channel (FC) physical interfaces and Fibre Channel Protocol (FCP) interfaces and services as per ANSI X3.230, X3.297 & X3.303.

5.10.6 CNA Interface Requirements

1. **[Conditional: DSC]** The system shall provide physical interfaces for Fibre Channel over Ethernet (FCoE) services over 10GbE physical interface in conformance with the ANSI T11 FC-BB-5 standard for FCoE with a Converged Network Adapter (CNA).

2. **[Conditional: DSC]** The system shall provide physical interfaces for Data Center Bridging (DCB), also known as Converged Enhanced Ethernet (CEE), features and functionality per the standards depicted below in Table 3.

DCB Standard	Description
IEEE 802.1Qbb for Priority-based Flow Control (PFC)	Per-Priority PAUSE adds fields to the standard PAUSE frame that allow a device to inhibit transmission of frames on certain priorities as opposed to inhibiting all frame transmissions.
IEEE 802.1Qaz for Enhanced Transmission Selection (ETS)	Enhanced Transmission Selection provides a means for network administrators to allocate link bandwidth to different priorities on the basis of a percentage of total link bandwidth.
IEEE 802.1Qaz Data Center Bridging Exchange Protocol (DCBX)	DCB Exchange is the mechanism in which peers can exchange capabilities to one another with Link Layer Discovery Protocol (LLDP).
IEEE 802.1Qau for Congestion Notification (CN)	Congestion Notification is a mechanism to transmit congestion information on an end-to-end basis per traffic flow.

Table 3 – Physical Interfaces for Data Center Bridging

5.10.7 IP Networking Requirements

1. **[Required: DSC]** The system shall meet the IPv6 requirements defined in Section 5.3.5 (in the DoD IPv6 Profile Section) for a simple server.

2. **[Conditional: DSC]** The system shall provide "Correspondent Node" functionality as per IETF RFC 3775 for "Mobility support in IPv6", section 9 for "Correspondent Node Operation". The system shall be a stationary node type that communicates with mobile computer node types. The Mobile IP functionality shall provide transport of HTTPS, SFTP, FTPS, SMTP, POP3, IMAP and SSHv2 protocols.

3. **[Required: DSC]** The system shall provide statically provisioned or dynamically adjusted large IP packet receive buffers for replication (mirroring) session traffic received on the Ethernet physical interfaces. The receive buffers may be statically provisioned or the operating system of the system may dynamically self-adjust the packet receive buffer size based upon measurements of the end-to-end path bandwidth, Maximum Segment Size (MSS), Round Trip Time (RTT) and the percentage of packet loss. The system shall provide a default and minimum IP packet receive buffer size of 2,048 KB per replication (mirroring) session. The system shall provide a statically provisioned or dynamically adjusting maximum IP packet receive buffer size of up to 8,192 KB per replication (mirroring) session.

These IP packet receive buffer size requirements are conceptually based upon either the Satellite or Transoceanic and Terrestrial Fiber Optic Cable end-to-end IP transport path models as depicted below in Table 4:

Path Model	Description
Transoceanic & Terrestrial Fiber Optic Cable	Where an end-to-end terrestrial OC-3 path with 155 Mbps of bandwidth that has a RTT of approximately 250 ms with packet loss of 0.01% or less. These characteristics are typical of a transoceanic and terrestrial fiber optic cable path between a pair of cities such as London and Tokyo. The 2,048 KB buffer size is suitable for these path characteristics.
Satellite	Where an end-to-end satellite DS1 path with 1.544 Mbps of bandwidth that has a RTT of approximately 600 ms with packet loss of 1.0% or greater. These characteristics are typical of a satellite path between two locations within the same VSAT footprint. The 8,192 KB buffer size is suitable for these path characteristics.

Table 4 – IP End-to-End Transport Path Models

4. **[Required: DSC]** The system shall provide an optimized congestion control (congestion avoidance) algorithm in TCP for avoidance of traffic loss on communications paths in high speed networks with high latency or large bandwidth-delay products.

NOTE: Two (2) examples of these algorithms that are currently implemented in modern operating systems are CUBIC TCP in Linux 2.6.19 and later, and Compound TCP (CTCP) in various Microsoft operating system products.

5.10.8 Name Services Requirements

1. **[Required: DSC]** The system shall provide Lightweight Directory Access Protocol (LDAP) directory services as per IETF RFC 4510.

2. **[Required: DSC]** The system shall provide Kerberos authentication service as per IETF RFC 4120.

3. **[Required: DSC]** The system shall provide Domain Name System (DNS) client functionality.

4. **[Required: DSC]** The system shall provide Domain Name Services (DNS) client side Load Balancing.

5. **[Required: DSC]** The system shall provide Network Information Service (NIS) client directory service functionality.

6. **[Required: DSC]** The system shall provide Network Information Service (NIS) Netgroups client directory service functionality.

7. **[Conditional: DSC]** The system shall provide NetBIOS over TCP/IP (NBT) Name Resolution & Windows Internet Name Service (WINS).

8. **[Required: DSC]** The system shall provide Internet Storage Name Service (iSNS) client functionality as per IETF RFC 4171.

9. **[Required: DSC]** The system shall provide Fibre Channel Name & Zone Service.

5.10.9 Security Services Requirements

1. **[Conditional: DSC]** The system shall provide IPsec as per RFC 4301 entitled “Security Architecture for the IP”.

2. **[Conditional: DSC]** The system shall provide Encapsulating Security Payload (ESP) as per RFC 4303 entitled “Encapsulating Security Payload (ESP)”.

3. **[Conditional: DSC]** The system shall provide Internet Key Exchange version 2 (IKEv2) as per RFC 4306 entitled “Internet Key Exchange version 2”.

4. **[Conditional: DSC]** The system shall provide a configurable Packet Filter (Firewall) service to block unauthorized access for intrusion prevention while permitting authorized communications. The Packet Filter service shall use a “stateless” design that does not degrade performance and shall filter all packets received based upon interface, source IP address, protocol, port, Type of Service (ToS) or Time To Live (TTL). The Packet Filter service shall provide a configuration policy for defining combinations of multiple packet match rules and processing actions.

5. **[Conditional: DSC]** The system shall work with the DoD HBSS software to protect data stored in the file systems in the attached disk array.

6.) **[Required: DSC]** The system shall provide encryption of data at rest at a minimum of AES-256 in accordance with FIPS 140-2 level 1 or higher to provide the following capabilities:

- Rapid crypto-shredding (destruction) of data in accordance with NIST 800-88, for tactical systems that operate in harm's way that may fall into enemy hands.
- Rapid recovery from sensitive data spills, where the wrong data is accidentally written to the wrong place.

7. **[Required: DSC]** The system shall comply with all appropriate STIGs to include the database STIG.

5.10.10 Interoperability Requirements

1. **[Required: DSC]** The system user interfaces, software, firmware, and hardware shall be compatible and interoperable with traffic transport and protection mechanisms of IP ASLAN and DISN WAN networks.

2. **[Required: DSC]** The system shall provide an Application Programming Interface (API) to enable interaction with other software and systems. The interactions shall include routines, data structures, object classes and protocols used to communicate between the consumer and implementer of the API. The API protocol and message format (e.g. XML) shall be subject to the specific vendor system operating system implementation.

5.10.11 Class of Service & Quality of Service Requirements

1. **[Required: DSC]** The system shall provide Class of Service (CoS) & Quality of Service (QoS) marking on egress traffic at layer 2 as per IEEE 802.1p and UCR2008 section 5.3.1.3.3 entitled "Class of Service Markings" and UCR section 5.3.1.3.4 entitled "Virtual LAN Capabilities". Traffic classification and marking must occur before the egress transmission of the Ethernet frame with a rule or policy engine that matches on various storage and management protocol types listed as follows and as offered by the system:

- NFSv3
- NFSv4
- NFSv4.1
- CIFSv1.0
- CIFSv2.0
- iSCSI
- FCoE

- HTTP/HTTPS/REST
- SFTP
- FTPS
- SSHv2
- SNMPv2
- SNMPv3
- User defined protocols (e.g. proprietary system to system mirroring protocol)

The marking is made in Ethernet VLAN tags by setting the priority value to between zero (0) and seven (7) inclusive for various traffic classes. These are to be used in the ASLAN, Non-ASLAN and extended networks for per-hop CoS & QoS traffic conditioning by the network elements.

2. **[Required: DSC]** The system shall provide Class of Service (CoS) & Quality of Service (QoS) marking on egress traffic at layer 3 as per UCR2008 section 5.3.3. Traffic classification and marking must occur before the egress transmission of the IP packet with a rule or policy engine that matches on various storage and management protocol types listed as follows:

- NFSv3
- NFSv4
- NFSv4.1
- CIFSv1.0
- CIFSv2.0
- iSCSI
- FCoE
- HTTP/HTTPS/REST
- FTPS
- SFTP
- SSHv2
- SNMPv2
- SNMPv3
- User defined protocols (e.g. proprietary system to system mirroring protocol)

IP packets are marked in the ToS field of the IPv6 packet header with Diff-serv Code Point (DSCP) values from zero (0) and sixty-three (63) inclusive. These are to be used in the ASLAN, Non-ASLAN and extended networks for per-hop CoS & QoS traffic conditioning by the network elements.

5.10.12 Virtualization Requirements

1. **[Required: DSC]** The system shall provide virtualized Data Storage Controller (vDSC) functionality and individual protocol server processes. The vDSC shall meet all of the requirements of a DSC with minor exceptions that are related to design and technical limitations

associated with the complete virtualization of an operating system. Examples include internal counters for attributes of the physical system, QoS traffic processing and per vDSC Mobile IP correspondent node binding cache limitations.

NOTE: Within the DSC system a vendor may integrate a third party component (s) that enables virtualization of heterogeneous file servers and provides a global namespace capability.

2. **[Required: DSC]** The vDSC capability within the system shall provide secure, Private Networking Domains (PND) for Ethernet, VLANs and IP that isolate the network domains of system units. The PND shall support the use of duplicate IP addresses and IP subnet address ranges among those of any other configured vDSC in the system. The PND shall provide a dedicated IP Forwarding Information Base (FIB) per vDSC.

3. **[Required: DSC]** The vDSC shall provide an individual CLI contexts with the full command set of the system with the scope of the commands limited to the individual vDSC CLI context.

4. **[Required: DSC]** The vDSC shall provide an programmatic API with the full command set of the system with the scope of the of the API commands limited to the individual vDSC context.

5. **[Required]** The vDSC capability within the system shall provide an individual Global Name Space (GNS) unique from the system or shall provide a Single Name Space that provides the capability to aggregate disparate hardware and storage architectures into a single file system. The GNS shall provide the capability to aggregate disparate and remote network based file systems, providing a consolidated view to reduce complexities of localized file management and administration. The GNS shall provide large working pools of disks, transparent data migration, and serve to reduce the number of storage mount points and shares. The Single Name Space shall be spread across multiple physical NAS heads all representing the same file system without replication. The Single Name Space shall include the ability to automatically tier data within the same file system.